

A Dynamic Axiomatic Approach to First-Price Auctions

Darrell Hoy*

Kamal Jain[†]

Christopher A. Wilkens[‡]

Abstract

The first-price auction is popular in practice for its simplicity and transparency. Moreover, its potential virtues grow in complex settings where incentive compatible auctions may generate little or no revenue. Unfortunately, the first-price auction is poorly understood in theory because equilibrium is not *a priori* a credible predictor of bidder behavior.

We take a dynamic approach to studying first-price auctions: rather than basing performance guarantees solely on static equilibria, we study the repeated setting and show that robust performance guarantees may be derived from simple axioms of bidder behavior. For example, as long as a loser raises her bid quickly, a standard first-price auction will generate at least as much revenue as a second-price auction.

We generalize this dynamic technique to complex pay-your-bid auction settings: as long as losers do not wait too long to raise bids, a first-price auction will reach an envy-free state that implies a strong lower-bound on revenue; as long as winners occasionally experiment by lowering their bids, the outcome will near the boundary of this envy-free set so bidders do not overpay; and when players with the largest payoffs are the least patient, bids converge to the egalitarian equilibrium. Significantly, bidders need only know whether they are winning or losing in order to implement such behavior.

Along the way, we find that the auctioneer’s choice of bidding language is critical when generalizing beyond the single-item setting, and we propose a specific construction called the *utility-target auction* that performs well. The utility-target auction includes a bidder’s final utility as an additional parameter, identifying the single dimension along which she wishes to compete. This auction is closely related to profit-target bidding in first-price and ascending proxy package auctions and gives strong revenue guarantees for a variety of complex auction environments. Of particular interest, the guaranteed existence of a pure-strategy equilibrium in the utility-target auction shows how Overture might have eliminated the cyclic behavior in their generalized first-price sponsored search auction if bidders could have placed more sophisticated bids.

1 Introduction

In 1961, Vickrey [28] initiated the formal study of auctions. He first considered common auctions of the day — including the first-price auction, the Dutch auction, and the English auction — and studied their equilibria. Vickrey observed that the English auction was, in theory, more robust because each player had a strategy that dominated all others regardless of other players’ bids.

*Northwestern University, Evanston, IL. Much of this work was done while the author was an intern at eBay Research Labs.

[†]eBay Research Labs, San Jose, CA.

[‡]University of California at Berkeley, Berkeley, CA. Much of this work was done while the author was an intern at eBay Research Labs. This work was supported in part by NSF award CCF-0964033

As a solution, he proposed¹ the second-price auction as a means to achieve the same robustness in a sealed-bid format. The subsequent development of auction theory largely followed Vickrey’s paradigm: existing auctions were evaluated in terms of their equilibria, meanwhile the field of mechanism design emerged with dominant strategy incentive compatibility as a *sine qua non*.

Fifty years later, it is apparent that Vickrey’s analysis does not always give best guide to implementing a real auction. In mechanisms without dominant strategies, Vickrey’s original concern still stands — equilibrium is a highly questionable predictor of outcome due (at least in part) to players’ informational limitations [28, 13]. Neither is dominant strategy incentive compatibility a magic solution: incentive compatible mechanisms have sufficiently many drawbacks that their real attractiveness rarely matches theory — the simple and elegant second-price auction has earned the title “Lovely but Lonely” [3] for its sparse use. Even the supposition that bidders will play strategies that are theoretically dominant is discredited by a wide variety of practical issues [18].

Dynamic analysis offers a powerful complement to Vickrey’s static approach. For example, certain behavior will be clearly irrational when an auction is repeated. Such reasoning was used by Edelman and Schwarz [10] in the context of the generalized second-price (GSP) ad auction — they analyzed a dynamic game to derive bounds on reasonable outcomes of the auction, then studied the static game under the assumption that these bounds were satisfied. Dynamic settings also introduce new pitfalls: Edelman and Ostrovsky [8] showed that the instability of Overture’s generalized first-price (GFP) ad auction could be attributed to its lack of a pure-strategy equilibrium.

We study repeated first-price auctions and show that they offer powerful performance guarantees. We begin with a static perspective and observe that the equilibrium properties of the auction depend significantly on the types of bids that bidders can express. We propose a generalization of the first-price auction called the *utility-target auction* that is closely related to profit-target bidding in first-price and ascending proxy package auctions [22, 7]. Like these package auctions, we show that the utility-target auction possess many advantages over incentive compatible mechanisms in a static equilibrium analysis, including revenue, simplicity, and transparency. More significantly, we show that *the same performance guarantees may be derived using only a few simple behavioral axioms and limited information* in a repeated setting. These dynamic results are particularly powerful because they do not require an a priori assumption that the auction will reach equilibrium — for example, assuming only that losers will not wait too long to raise their bids *the auctioneer’s revenue satisfies a natural lower bound regardless of whether bidders’ behavior converges to equilibrium*. Moreover, bidders need only know if they are winning or losing to implement the dynamics. We build on these axioms to demonstrate behavior that offers progressively stronger performance guarantees, culminating with a set of axioms that together imply convergence to the egalitarian equilibrium.

First-Price Auctions. The virtues of a first-price auction — and other auctions in the pay-your-bid family — arise from its simplicity. From the bidders’ perspective, the pay-your-bid property offers transparency, credibility, and privacy: not only is the auction easy to understand, but it ensures that the auctioneer cannot cheat (say, by unreasonably inflating the reserve price in a repeated auction) and allows a bidder to participate without expressing her true willingness to pay.

The auctioneer can also benefit from this simplicity because players’ bids represent guaranteed revenue. By comparison, the revenue from a dominant strategy incentive compatible auction is

¹While Vickrey was the first to discover the second-price auction in the economics literature, it has been used in practice as early as 1893 [20].

almost always less than the bids and, in the most general settings, may even be zero [3, 25]. Supposing a first-price auction reaches equilibrium, a variety of work presents settings where they generate more revenue for the seller than their incentive compatible brethren [22, 19, 17] (though they may also generate less revenue [21]).

Yet, running a first-price auction is risky. While first-price auctions have been quite successful in settings like treasury bill and procurement auctions, Overture’s generalized first-price (GFP) auction for sponsored search advertising was erratic: bids rapidly rose and fell in a sawtooth pattern, rendering the auction unpredictable and depressing revenue [8]. As a result, the sponsored search industry has moved to a generalized second-price (GSP) auction that leverages the intuition of the second-price auction to disincentivize small adjustments to a player’s bid.

The challenges of a first-price auction are many and complex. Vickrey identified a major source of risk in the first-price single-item auction: since a rational bidder’s optimal bid depends on other players’ bids, actual behavior will depend on beliefs about others’ strategies. A first-price auction also requires bidders to strategize, a task that is may be difficult and expensive. At best, players will be in a Bayesian equilibrium, and, at worst, they will be completely unpredictable. Indeed, predicting the outcome of a first-price auction lies at the center of a lively debate between experimental and theoretical economists [13].

Experience with GFP highlights another potential pitfall of first-price auctions: when generalized beyond the single-item setting, a first-price auction may not have a pure-strategy equilibrium. Edelman and Ostrovsky [8] showed that this was the case with GFP and demonstrated how it generated the rapid sawtooth behavior seen in practice. Our goal is to demonstrate that how proper design coupled with dynamic arguments can support strong performance guarantees.

The Utility-Target Auction. The equilibria and performance of a pay-your-bid auction depend on its implementation. Within the pay-your-bid constraint, the auctioneer chooses the form of players’ bids, potentially restricting or broadening the bids that players may express.

The historical performance of the GFP ad auction exemplifies the importance of choosing a good bidding language. In the GFP auction, advertisers placed a single bid and paid the bid price for each click regardless of where their ads were shown. In retrospect, the rapid sawtooth motion observed in bids is not surprising because the auction had no pure-strategy equilibrium [9, 8]; however, we show that a pure-strategy equilibrium would have existed if bidders could have placed more expressive bids, such as bidding different prices for each slot.

A natural question arises: what are good bidding languages and how complicated must a language be to offer good performance? In GFP, the bidding language is precisely sufficient to represent any possible valuation function; hence, it is possible that bids may need to be more expressive than the space of valuation functions.

We show that the overhead required for a good bidding language is at most a single value: it is sufficient to ask bidders for their valuation function and their final desired utility. We call such an auction a *utility-target auction*: a player’s bid includes a specification of her value for every outcome and a single number representing the utility-target that she requests regardless of the outcome. Her payment is her claimed value for the final outcome minus the utility-target that she requested, and the auctioneer chooses the outcome that maximizes the total payment. In essence, the utility-target auction isolates the single dimension (utility) along which a bidder truly wishes to strategize.

We begin with a static analysis of the utility-target auction’s equilibria. We first show that the

utility-target auction is quasi-incentive compatible: a bidder never has an incentive to misreport her valuation function — it is always sufficient for her to manipulate the utility-target she requests. Moreover, we show that a pure-strategy equilibrium always exists and that the egalitarian equilibrium is efficiently computable. These results are closely related to profit-target equilibria in package auctions [22].

Next, we show that the utility-target auction offers good equilibrium performance. Similar to the approach of Edelman, Ostrovsky, and Schwarz [9] on the generalized second-price (GSP) auction, we show that all equilibria satisfying a natural envy-free criterion have good performance. First, such equilibria are efficient and generate at least as much revenue as the incentive compatible Vickrey-Clarke-Groves (VCG) mechanism. Moreover, they generate revenue even when the incentive compatible mechanisms fail — the revenue of the envy-free equilibria of a utility-target auction all meet an intuitive benchmark we call the second-price threat, even settings where the VCG mechanism may make little or no revenue. Again, this bound is related to the core property of profit-target equilibria in package auctions [22, 7].

Dynamic Analysis through Behavioral Axioms. A significant novelty of our work is our use of simple behavioral axioms to prove guarantees on the performance of utility-target auctions.

Dynamic arguments are generally fraught with peril: in addition to being difficult to prove, more complex auctions (or markets, or games) require more complex bidding behavior to converge to an equilibrium and therefore sacrifice robustness. For example, Walrasian tâtonnement² [29] is perhaps the earliest concrete dynamic procedure proposed in economics — it converges in general markets when modeled as a particular continuous process [27, 1] but may or may not converge as a discrete process [4, 6]. More recent results have sought stronger guarantees, e.g. by showing that players’ behavior will converge to equilibrium in repeated games as long as their learning strategies are “adaptive and sophisticated” [24] or no-regret [14, 12]. However, these properties are sufficiently complicated that it is difficult to evaluate whether players’ strategies indeed satisfy them in practice.

In contrast, we build simple behavioral axioms and use them to prove performance guarantees. Our first axioms are that (a) a bidder who is losing will raise her bid to try to win, and (b) a bidder who is losing is more impatient than a bidder who is winning. After formalizing these axioms in the context of utility-target auctions, we show that the auction will eventually reach an outcome that satisfies a natural notion of envy-freeness and, by extension, a natural second-price type bound on revenue. Significantly, this result neither implies nor requires that players’ bids converge to a steady-state. Moreover, bidder behavior requires only knowing whether one is winning or losing, not the precise bids of other players.

Next, we show that bidders will not overpay if two more axioms are also satisfied, namely that (c) bidders who are winning will try to lower their bid to save money. Axioms (a)-(c) guarantee that bids will ultimately remain close to the boundary between envy-free and non-envy-free outcomes, a boundary which contains the envy-free equilibria. These axioms offer a degree of robustness, since bids will seek this boundary even as bids and ads change.

Finally, we show that bids will converge to the egalitarian equilibrium — the equilibrium that distributes utility most evenly — if a fourth axiom is satisfied. The fourth axiom concerns the timing of raised bids: (d) the bidder who has the most value at risk is the least patient and

²To justify market equilibrium as a predictor of actual market behavior, Leon Walras described a dynamic procedure called tâtonnement that might converge to it.

therefore raises her bid first. When bidder behavior satisfies all five axioms (a)-(d), we show that bids will converge to the egalitarian equilibrium. Together, these results offer powerful guarantees about the performance of a utility-target auction in a repeated setting.

Related Work. Our utility-target auction is most closely related to first-price package auctions [5] and the ascending proxy auction [22]. Profit-target bidding in these auctions is closely related to quasi-truthful bidding in utility-target auctions, and the static properties we prove in Section 4 all have direct analogues. In contrast, the utility-target auction can be applied beyond the package auction setting (e.g. to ad auctions), and our dynamic analysis is entirely new, a more general confirmation of Milgrom’s postulate that profit-target equilibria “may describe a central tendency for some kinds of environments” [22].

Auctions in which a player’s bid directly specifies her payment are known as *pay-your-bid* auctions. The first-price auction, as well as the Dutch and English auctions, are members of this family. Our utility-target auction is closely related to first-price and ascending proxy package auctions [22]. Engelbrecht-Wiggans and Kahn [11] explored multi-unit, sealed-bid pay-your-bid auctions and found their equilibria to be substantially different from the standard first-price auction — the core issue they encounter is the same one arising in GFP.

A key reason repeated auctions may admit more robust performance guarantees is that bidders can learn about others’ valuations. A similar informational exchange is present in and a motivation for classic ascending auctions. In addition to his discussion of ascending proxy auctions [22], Milgrom offers a broad discussion of this literature in [23]. Some recent work studies ascending auctions for position auctions like sponsored search [9, 2].

Our work can also be seen through the lens of *simple versus optimal* mechanisms [16]. The general goal of this line of research is to design a mechanism that is simple and transparent while (possibly) sacrificing efficiency or revenue. For example, Hart and Nisan analyze the tradeoff between the number of different bundles offered to a buyer and an auction’s performance [15]. By comparison, our results show that a first-price auction can guarantee good performance when the bidding complexity is only slightly larger than that of the valuation functions.

2 Definitions and Preliminaries

The *utility-target auction* is a generalization of the first-price auction. Its key feature is an extra utility-target parameter in the bid — this parameter highlights the key dimension along which bidders care to compete. It gives bidders sufficient flexibility to guarantee the existence of pure-strategy equilibria while minimizing the communication required between the bidders and the auctioneer.

2.1 First-Price and Pay-Your-Bid Auctions.

An auction is a protocol through which players bid to select an outcome. A standard sealed-bid auction can be decomposed into three stages: (1) each player i submits a bid b_i , (2) the auctioneer uses players’ bids to pick an outcome o from a set \mathcal{O} , and finally (3) each player i pays a price p_i . The final utility of player i is given by $v_i(o) - p_i$, where $v_i(o) \geq 0$ denotes i ’s value for the outcome o , i.e. $v_i \in V_i$ is i ’s valuation function (drawn from a publicly known set V_i).

From this perspective, the standard first-price auction is described as follows: (1) each player submits a single number $b_i \in \mathbb{R}$, (2) the auctioneer chooses to give the item to the player i^* who submits the largest bid b_i , and (3) the winner i^* pays b_{i^*} and everyone else pays zero. For comparison, the second-price auction is identical to the first-price auction except that the price paid is equal to the second-highest value of b_i .

When the outcomes are few, we will use v_i and b_i to denote the profile of values and bids across outcomes, e.g., $v_i = (1, 1.5, 0)$.

When considering settings beyond the single-item auction there are many ways to generalize the first-price auction. Even within the single-item setting, the auctioneer could choose an arbitrary encoding for players' bids. Moreover, the auctioneer might choose an encoding that changes the space of possible bids, e.g by forcing bidders to place integer bids when values are actually real numbers. In such cases, the principle feature that we wish to preserve is that the winner "pays what she bid," or alternatively that a player's bid precisely specifies her payment. Formally, we say that such an auction has the pay-your-bid property:

Definition 1 *An auction has the pay-your-bid property if the payment p_i depends only on the outcome o and i 's bid b_i (it does not directly depend on others' bids).*

The first-price auction as described above clearly satisfies this property while a second-price auction does not.

Not all sealed-bid pay-your-bid auctions are equivalent. Edelman et al. [9] showed that GFP, where the set of possible bids is precisely V_i , did not have a pure-strategy equilibrium:

Observation 1 *The pay-your-bid property does not guarantee the existence of a pure-strategy equilibrium in a sealed-bid auction when the space of bids is the same as the space of valuation functions.*

Moreover, as we discuss in Section 5, any pay-your-bid ad auction where bids are restricted to a subset of V_i must suffer in terms of its welfare and revenue guarantees. Thus, it is important to consider auctions that allows bids $b_i \notin V_i$. This motivates us to introduce the utility-target auction, a sealed-bid pay-your-bid auction that allows such bids and always has pure-strategy equilibria with strong performance guarantees.

2.2 Utility-Target Auctions

A utility-target auction is a sealed-bid pay-your-bid auction with a special bidding language. A player's bid specifies payments using two pieces of information: her valuation function and the amount of utility she requests (a single real number). Her payment for an outcome is her (claimed) valuation for that outcome minus the utility that she specified in her bid. Formally:

Definition 2 *A utility-target auction for a finite outcome space \mathcal{O} is defined as follows:*

- *A bid is a tuple $b_i = (x_i, \pi_i)$ where $x_i \in V_i$ is a function mapping outcomes $o \in \mathcal{O}$ to nonnegative values and π_i is a real number. We call the parameters x_i and π_i the value bid and utility-target bid respectively.*
- *A bidder's effective bid for outcome o is*

$$b_i(o) = \max(x_i(o) - \pi_i, 0) .$$

Note this may generate $b_i \notin V_i$ when the set V_i is sufficiently restricted.

ALGORITHM 1: A generic utility-target auction.

input : Players' bids $b_i = (x_i, \pi_i)$
output: An outcome o^* and first-price payments p_i .

- 1 Let $b_i(o) = \max(0, x_i(o) - \pi_i)$; // $b_i(o)$ is i 's effective bid for outcome o .
- 2 Compute $o^* = \operatorname{argmax}_o \sum_{i \in [n]} b_i(o)$; // Choose the outcome with the highest total bid.
- 3 For all i , set $p_i = b_i(o^*)$; // Each player pays what she bid.

- The auctioneer chooses the outcome $o^* \in \mathcal{O}$ that maximizes $\sum_{i \in [n]} b_i(o)$. Ties are broken in favor of the most-recent winning outcome when applicable.
- When the outcome is o , bidder i pays $p_i(o) = b_i(o)$ and derives utility $u_i(o) = v_i(o) - b_i(o)$. Note that if a bidder reports $x_i = v_i$, then $u_i(o) = \pi_i$ whenever $v_i(o) \geq \pi_i$.

A generic utility-target auction is illustrated in Algorithm 1.

3 Quasi-Truthful Bidding

An idealist's intuition for the utility-target auction is that players truthfully reveal their valuation function through their value bids (i.e. they bid $x_i = v_i$) and then use the utility-target bid π_i to strategize. Clearly, bidders need not follow this ideal; however, it turns out that they have no incentive to do otherwise — the utility-target auction is quasi-truthful in the sense that for any bid a player might consider, there is another bid in which she reveals v_i truthfully and obtains at least as much utility:

Lemma 3.1 (Quasi-Truthfulness) *Fix the total bid of players $j \neq i$ for all outcomes, i.e. fix $\sum_{j \in [n] \setminus \{i\}} b_j(o)$ for all o , and suppose ties are broken according to a fixed total-ordering on outcomes. If bidder i gets u_i^I by bidding (x_i^I, π_i^I) , then she gets the same utility u_i^I by bidding (v_i, u_i^I) .*

Significantly, this implies bidder i always has a quasi-truthful best-response.

Proof: Since ties are broken according to a fixed total ordering, the outcome is fully specified by the total bids for each outcome (i.e. by $\sum_{i \in [n]} b_i(o)$ for all o). Thus, given $\sum_{j \in [n] \setminus \{i\}} b_j(o)$ and a bid $b_i^I = (x_i^I, \pi_i^I)$ for i , the outcome o^I is uniquely defined. Let π_i^I be the utility i gets by bidding (x_i^I, π_i^I) , i.e.

$$u_i^I = v_i(o^I) - b_i(o^I) = v_i(o^I) - \max(x_i^I(o^I) - \pi_i^I, 0) .$$

Now suppose i bids $b_i^Q = (v_i, u_i^I)$ instead of (x_i^I, π_i^I) . There are two possible results of this change:

- *The outcome doesn't change.* If the outcome doesn't change, then i gets the same utility by construction.
- *The outcome changes to $o^Q \neq o^I$.* Notice that i did not change the amount bid for outcome o^I , so the total bid for o^I did not change. Given this and the tie-breaking rule, the only way the outcome can switch from o^I to o^Q is if the total bid for o^Q strictly increased. Given that $\sum_{j \in [n] \setminus \{i\}} b_j(o^Q)$ is fixed, this implies i 's bid for o^Q increased, i.e. $b_i^Q(o^Q) > b_i^I(o^Q) \geq 0$.

Next, by definition of a utility-target auction, $b_i(o) > 0$ implies $x_i(o) > \pi_i(o)$. Since $b_i^Q(o^Q) \geq 0$, this implies $v_i(o^Q) > \pi_i^Q$, from which it immediately follows that i 's final utility in o^Q will be π_i^Q .

In either case, i 's final utility is precisely u_i^I , so i is indifferent between bidding (x_i^I, π_i^I) and (v_i, u_i^I) . ■

4 Static Equilibrium Analysis

We begin by studying the utility-target auction from a static perspective and show that they offer strong revenue and welfare guarantees. First, we show that pure-strategy equilibria always exist:

Theorem 4.1 *A utility-target auction with n outcomes always has a pure-strategy cooperatively envy-free (defined below) equilibrium that is computable in time $\text{poly}(n)$.*

Specifically, the egalitarian equilibrium exists and is efficiently computable by Algorithm 2 (proof omitted).

Next, we show that such cooperatively envy-free equilibria not only maximize welfare but offer as much revenue as the VCG mechanism as well as a new revenue benchmark we call the second-price threat (defined below):

Theorem 4.2 *Any cooperatively envy-free equilibrium of a utility-target auction*

1. *maximizes social welfare,*
2. *dominates the revenue of the VCG mechanism,*
3. *and has revenue lower-bounded by the second-price threat.*

We formalize and prove the theorem below.

4.1 Cooperatively Envy-Free Equilibria

While a typical utility-target auction may have many equilibria, some of them are unrealistic in repeated auctions. In particular, it is possible to have an equilibrium in which a group of “losers” envy the “winners” — the losers would be happy to collectively raise their bids to make an alternate outcome win, but the outcome is an equilibrium because no single bidder is willing to raise her bid high enough. In a repeated setting, one would expect all the losers to eventually raise their bids.

For example, consider a setting with three bidders (A, B, C) and three outcomes $(1, 2, 3)$, in which the first two bidders are symmetric and value the first two outcomes and the third bidder values only the third outcome. Let the specific values, indexed by outcome, be

$$\begin{aligned} v_A &= (1, 1.5, 0), \\ v_B &= (1, 1.5, 0), \\ v_D &= (0, 0, 2). \end{aligned}$$

Now, let A and B bid for the first outcome, and C bid for the third outcome, with bids: $b_A = (1, 0, 0)$, $b_B = (1, 0, 0)$ and $b_C = (0, 0, 2)$.

Bidders A and B would prefer the second outcome to the first as they see a value of 1.5 instead of 1. Moreover, they would be happy to make the second outcome win by cooperating and each bidding $1 + \epsilon$. However, since both are bidding 0 for the second outcome, neither can unilaterally cause the second outcome to win, making this outcome an equilibrium. The problem in this example is that, at a total price of 2, bidders A and B would prefer that the second outcome wins. In a sense, bidders A and B in the second outcome envy the deal they received in the first outcome.

Hence, we are interested in bids such that players have no incentive to cooperatively deviate to get a better outcome. We will call such a set of bids *cooperatively-envy free*.

To define such a notion, we must also consider bidders who are happy with the winning outcome. Consider a four bidder setting with three possible outcomes, with the following values:

$$\begin{aligned} v_A &= (1, 1.5, 0), \\ v_B &= (1, 1.5, 0), \\ v_C &= (1, 0.5, 0), \\ v_D &= (0, 0, 2). \end{aligned}$$

In this case, C cannot get a better deal from the second outcome, so she will not cooperate with A and B . In order to win, A and B must collectively bid 1.75 in the second outcome to make up the deficit between it and the winning outcome (which they are willing to do).

Definition 3 *The set of bids $\{b_i\}_{i \in [n]}$ are cooperatively envy-free (CEF) if there is no subset of bidders $\mathcal{B} \subseteq [n]$ who would prefer to cooperatively pay the extra money required to make an alternate outcome o win over the current winner o^* .*

Formally, a set of bids is cooperatively envy-free if

$$\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) \leq \sum_{i \in [n]} b_i(o^*) - b_i(o)$$

for all outcomes o .

The CEF constraints are similar to the core property described by Milgrom [22] in package auctions (as well as notions like group-strategyproofness); however, the notion of a CEF outcome is weaker. For example, it does not require that bidders are playing equilibrium strategies.

Equilibria that are CEF have nice properties analogous to those of core equilibria in package auctions. The following claims are straightforward and proven in the appendix of the full version of the paper:

Claim 1 *CEF bids maximize welfare.*

Claim 2 *The revenue from CEF bids dominate that of the VCG mechanism: every player pays at least as much in the CEF equilibrium as she would in the VCG mechanism.*

ALGORITHM 2: An algorithm for computing the egalitarian equilibrium in a utility-target auction.

input : A utility-target auction problem.

output: The egalitarian equilibrium bids $b_i^* = (v_i, \pi_i^*)$.

- 1 Set all bids to $(v_i, 0)$. Call the socially optimal outcome o^* .
 - 2 Increase π_i for all bidders uniformly until some bidder i reaches $\pi_i = v_i(o^*)$ or a CEF constraint would be violated for some outcome o .
 - 3 Fix the bids of the newly-constrained advertisers.
 - 4 Repeat (2) and (3), lowering only unfixed bids until all bidders are fixed.
-

4.2 The “Second-Price Threat”

The revenue of a CEF equilibrium also meets or exceeds a benchmark we call the second-price threat. The revenue of the second-price auction has a convenient intuition: the price paid by the winner should be at least as large as the maximum willingness to pay of any other bidder. We can ask the same question in more general settings: how much would “losers” be willing to pay to get an outcome o instead of the socially optimal outcome o^* ? In general, player i should be willing to pay up to $v_i(o) - v_i(o^*)$ to help o beat o^* , hence we can generalize the intuition of the second-price auction to give a natural lower bound on the revenue the auctioneer might hope to earn:

Definition 4 *The second-price threat for outcome o^* is given by*

$$\max_{o \in \mathcal{O}} \sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) .$$

This bound is particularly powerful in cases where bidders share value for an outcome (cases where VCG would make little or no revenue). For example, consider the following 4-bidder, 2-outcome setting:

$$v_A = (1, 0)$$

$$v_B = (1, 0)$$

$$v_C = (1, 0)$$

$$v_D = (0, 2)$$

In a VCG auction, nobody pays anything. However, a naïve auctioneer might expect the first outcome to win, with A , B , and C paying a total of \$2 (the second-price threat) since they are beating D .

The CEF constraints quickly imply that a CEF outcome generates at least as much revenue as the second-price threat:

Claim 3 *The revenue in any CEF outcome is lower-bounded by the second-price threat.*

Proof is given in the appendix of the full version of the paper.

5 Utility-Target Auctions for Sponsored Search

Sponsored search advertising demonstrates the benefits of a utility-target auction. The standard auction in this setting is the generalized second-price (GSP) auction; however, it (and incentive-compatible VCG mechanisms) lack transparency: payments are complicated to compute and bidders must trust the auctioneer not to abuse their knowledge when an auction is repeated. Moreover, its performance may degrade when using more accurate models of user behavior [26] and advertiser value [17]. It can have misaligned incentives when parameters are estimated incorrectly [30]. Some of these problems would be solved by a first-price auction; however, Overture’s implementation of GFP demonstrated that such schemes might be highly unstable. A utility-target auction offers the benefits of a pay-your-bid auction without the instability of GFP.

5.1 The Utility-Target Ad Auction

We illustrate a utility-target auction in the standard model of sponsored search: n advertisers compete for $m \leq n$ slots associated with a fixed keyword. An advertiser’s value depends on the likelihood of a click, called the click-through-rate (CTR) c , and the value v to the advertiser of a user who clicks. The CTR c is separable into a parameters β_i that depends on the advertiser and α_j that depends on the slot, so the expected value to advertiser i for having her ad shown in slot j is $c_{i,j}v_i = \alpha_j\beta_i v_i$. As is standard, we assume that slots are naturally ordered from best ($j = 1$) to worst ($j = m$), i.e. $\alpha_j \geq \alpha_{j'}$ for all $j < j'$. Without loss of generality, we assume bidders are ordered in decreasing order of bid, i.e. $b_1 \geq b_2 \geq \dots \geq b_n$.

The auctioneer chooses a matching of advertisements to slots and charges an advertiser a per-click price ppc_i . For example, in the GFP auction, advertisers submitted bids b_i representing their per-click payment and paid $ppc_i = b_i$ whenever a their ads were clicked. Similarly, in the standard GSP auction, bidder i is charged according bid of the next highest bidder.³ To account for differences in CTRs, this quantity is normalized by β so that bidder i pays a per-click price of $ppc_i = \frac{\beta_i}{\beta_{i+1}}b_{i+1}$.

In a *utility-target auction*, bidders submit both their per-click value x_i and the utility-target bid π_i (the utility that they request). The auctioneer picks the assignment $j(i)$ maximizing

$$\sum_{i \in [n]} \max(0, \alpha_{j(i)}\beta_i x_i - \pi_i)$$

and charges i so that her expected payment is

$$\mathbf{E}[p_i] = \max(0, \alpha_{j(i)}\beta_i x_i - \pi_i) \ .$$

There are at least two interesting ways the utility-target auction can be implemented. The first implementation charges

$$ppc_i = \max\left(0, x_i - \frac{\pi_i}{\alpha_{j(i)}\beta_i}\right)$$

to achieve the desired expected payment. In effect, it uses the utility request π_i to compute a different per-click bid for each slot. A practical downside to this implementation is that the

³The designers of the GSP auction intended it to inherit the incentive compatibility of the second-price auction. It does not; however, it has the nice property that bidder i pays the minimum amount required to win the slot that she received.

payments are still somewhat complicated from the bidders’ perspectives; however, the auctioneer could mitigate this problem by publishing CTRs and displaying the per-click payments in the bidding interface.

An alternative implementation of the utility-target auction pays a rebate of π_i regardless of whether a click occurred and charges precisely $ppc_i = x_i$ when a click occurs. This auction is even simpler from the bidders’ perspective; however, when a click does not occur the auctioneer will be paying the bidder (in expectation the bidder still pays the auctioneer). This implementation of the utility-target auction is illustrated in Algorithm 3.

Such a utility-target auction offers many benefits over existing auction designs like GSP and VCG. As noted earlier, a first-price auction directly increases transparency and simplicity from the bidders’ perspective. Even if bidders reveal their true valuation functions v_i , the pay-your-bid property ensures that increasing a reserve price will not increase payments unless bidders subsequently raise their bids.

The auction also easily generalizes to more complicated bidding languages. Whereas the welfare and revenue performance of GSP degrades (albeit gracefully) [26] when considering externalities imposed by the presence of competing ads, the reasonable (CEF) equilibria of the utility-target auction guarantee good performance. An open question of [17] is to find an auction that performs well in multi-slot settings when multiple bidders can benefit from clicks on the same ad (e.g. Microsoft and Samsung both benefit from an ad for a Samsung laptop running Windows) — the utility-target auction offers good revenue guarantees in these ‘coopetitive’ ad auctions with multiple slots. The utility-target auction is also less sensitive to estimation errors in the CTRs. As shown in [30], incentive-compatibility can be broken because the auctioneer only knows estimates of the α and β parameters. Informally, the utility-target auction is much less sensitive to such errors because the payments need not explicitly depend on the auctioneer’s estimates.

5.2 Utility-Target vs. GFP

Juxtaposing GFP with the utility-target auction illustrates the benefits of a more complex bidding language. GFP is identical to the utility-target ad auction except that bids contain only the per-click payment x_i and not the utility-target bid π_i . Consequently, a player’s bid necessarily offers the same per-click payment regardless of the slot won by the bidder. By comparison, the utility-target auction permits bids that encode a different per-click payment depending on the slot in which an ad is shown.

In retrospect, it is easy to see that different per-click bids are important for a good pure-strategy equilibrium. In GFP, all advertisers who are shown must bid so that $\beta_i x_i$ is the same, otherwise some bidder can lower her value of x_i without changing her assignment; however, if this is true, then some bidder can move up to the top slot by bidding $x_i + \epsilon$. In fact, any bidding language that requires the same per-click payment for all slots could not have a pure-strategy equilibrium unless the potential benefit of being in the top slot was less than the effective bid increment required to get there. This necessarily weakens any revenue guarantees and, worse, implies that the auction cannot differentiate between the winning bidders to pick the best ordering of ads.

As noted earlier, the existence of a pure-strategy equilibrium is directly related to the dynamic performance of the auction. Edelman and Ostrovsky [8] discuss how the lack of such an equilibrium naturally leads to sawtooth cycling behavior in GFP, as bidders alternate between increasing their bids to compete for higher slots and decreasing their bids to avoid overpaying for the slots they have. They also show that this cyclic behavior potentially reduced revenue below that of the

ALGORITHM 3: Autility-target auction for search advertising.

input : Bids $b_i = (x_i, \pi_i)$.
output: An assignment of advertisements to slots, per-click payments ppc_i , and unconditional payments r_i .

- 1 For each bidder i and slot j , compute $\mathbf{E}[p_{i,j}] = \max(\alpha_j \beta_i x_i - \pi_i, 0)$;
- 2 Compute assignment $j(i)$ of advertisements to slots that maximizes $\sum_{i \in [n]} \mathbf{E}[p_{i,j(i)}]$;
- 3 **if** $\mathbf{E}[p_{i,j(i)}] > 0$ **then**
 - | Always pay i the rebate $r_i = \pi_i$;
 - | Whenever i 's ad is clicked, charge $ppc_i = x_i$;
- else**
 - | Do not charge/pay anything to i ;

VCG mechanism. In contrast, Theorem 4.1 shows that utility-target auctions have pure-strategy equilibria, and Theorem 4.2 shows that revenue at equilibrium dominates the VCG mechanism; moreover, our dynamic results show that bids will naturally approach this equilibrium (or the set of such equilibria) as bidders adjust their utility targets.

6 Dynamic Analysis

In this section, we consider the behavior of utility-target auctions in a very simple dynamic setting, and under very simple assumptions. We show that a few rules and simple knowledge of whether one is winning or losing are enough to guarantee the revenue and welfare bounds from Theorem 4.2.

Following Lemma 3.1, we assume that bidders are quasi-truthful and report bids of the form (v_i, π_i) . Bidders compete using the utility-target terms π_i and employ strategies to optimize their utility.

Winners and Losers. Our dynamic axioms are based on a natural decomposition of bidders into winners and losers. In a standard first-price auction, the winner is the bidder who gets what he wants — the item — and the losers are those who do not get what they want. In a utility-target auction, a bidder effectively reports her valuation v_i and requests that the auctioneer give her a certain utility π_i . This suggests partitioning bidders into winners and losers based on whether a bidder gets the utility she requests, giving us the following formal definition:

Definition 5 (Winners and Losers) *A winner is a bidder who gets the utility she requests, i.e. if o^b is the outcome of the auction, then i is a winner if and only if*

$$u_i(o^b) = v_i(o^b) - b_i(o^b) = \pi_i \text{ .}$$

Any bidder who is not a winner is a loser.

Observation 2 *Bidder i is a winner if and only if $v_i(o^b) \geq \pi_i$ and is always a winner when $\pi_i = 0$. When bidder i is a loser, $u_i(o^b) < \pi_i$.*

Note that this definition does not coincide with the standard definition of winners and losers in a single item auction because a bidder who does not get the item is still a winner if $\pi_i = 0$.

Raising and Lowering Bids. In a dynamic setting, we want to think about how winners and losers manipulate π_i . In the utility-target auction, the effective bid b_i (what bidder i is actually offering to pay) and the utility-target term π_i move in opposite directions, so when we talk about raising i 's bid we are talking about decreasing the utility-target term π_i :

Definition 6 (Raising and Lowering Bids) *We say that bidder i raises her bid from (v_i, π_i) if she chooses a new bid (v_i, π'_i) where $\pi'_i < \pi_i$, i.e. she raises her bid if she decreases her utility-target bid.*

Similarly, a bidder who lowers her bid correspondingly increases her utility-target bid from π_i to $\pi'_i > \pi_i$.

Importantly, our definition of winners and losers shares a natural property with the standard definition: winners cannot benefit by offering to pay more, and losers cannot benefit by offering to pay less:

Claim 4 *Fixing other players' bids, a loser cannot increase her utility by raising her bid. Likewise, a winner cannot increase her utility by lowering her bid.*

The claim is straightforward to prove.

Our definition of winners and losers also has a new property that is important:

Claim 5 *A loser can always raise her bid in a way that weakly increases her utility.*

Proof: Suppose i is a loser bidding (v_i, π_i) and receiving utility $u_i < p_i$. If she raises her bid to (v_i, u_i) , Lemma 3.1 says that she will receive utility of precisely u_i , making her a winner. ■

In our model, bidders locally adjust their bids by ϵ . To mimic settings where auctions happen frequently and no two bidders move simultaneously, bid changes are modeled as asynchronous events. As noted earlier, our model assumes players bid quasi-truthfully, that is, they always submit their true valuation functions in their bids. As a result, the history of the auction is characterized by a sequence of utility-target vectors π^0, \dots .

We assume that $0 \leq \inf_o v_i(o)$ and $\sup v_i(o) < \infty$, so utility-targets will always lie in the finite interval $[0, \sup v_i(o)]$. Unless a player's utility-target hits the boundary of this interval, all bid changes are made in increments of ϵ .

Notions of Convergence. We will show that progressively stronger assumptions imply progressively stronger convergence guarantees. Our first results show that bids will eventually be close to the set of CEF (or non-CEF) bids. As noted earlier, the utility-targets π are sufficient to characterize bidders' strategies, so we define \mathcal{C} to be the set of all such utility-target:

Definition 7 (The CEF Set) \mathcal{C} is the set of all utility-target vectors π where the quasi-truthful bids (v_i, π_i) produce a cooperatively envy-free outcome.⁴

The set $\overline{\mathcal{C}}$ is the set of all utility-target vectors which are not CEF, i.e. $\overline{\mathcal{C}} = \mathbb{R}_+^n \setminus \mathcal{C}$.

⁴Note that membership in \mathcal{C} depends on both the vector π and the outcome chosen by the auction. This is because certain utility-target vectors π will be in \mathcal{C} if ties are broken in favor of o^* but not if ties are broken in favor of a suboptimal outcome.

Significantly, \mathcal{C} is never empty. In particular, it always contains the 0 vector ($0^n \in \mathcal{C}$).

Since bidders are continually experimenting with their bids, it is not realistic to expect bids to explicitly converge to \mathcal{C} ; rather, they will remain close. For a set of bids π , let π_ϵ denote the set of bids that are close to some vector in π , i.e.

Definition 8 Let S_ϵ be the set of all utility-targets π which are close to some vector in π coordinate-wise. Formally,

$$S_\epsilon = \{\pi \mid \exists \pi' \in S \text{ s.t. } \|\pi - \pi'\|_\infty \leq \epsilon\} .$$

In particular, we will care about the sets \mathcal{C}_ϵ and $\bar{\mathcal{C}}_\epsilon$, the sets representing bids close to being CEF and close to being not CEF, respectively.

Next we define the convergence of an auction to utility-target bids π :

Definition 9 An auction converges to a set of utility-targets S if, for any $\delta > 0$, there exists a sufficiently small bid adjustment parameter ϵ for which the auction always reaches a utility-target π such that all future bids are in S_δ .

Our strongest result will show that bids converge to the egalitarian equilibrium:

Definition 10 The egalitarian equilibrium is the CEF equilibrium which distributes utility as evenly as possible. Formally, for each equilibrium let u_\uparrow be the vector of bidders' utilities with its coordinates sorted in increasing order. The egalitarian equilibrium is the one for which u_\uparrow is lexicographically maximized.

An auction converges to the egalitarian equilibrium π^E if it converges to $\{\pi^E\}$.

6.1 Axioms and Results

Our convergence theorems show that progressively stronger assumptions about bidder behavior lead to progressively stronger convergence results.

Our first axiom of bidder behavior captures some intuition about how winners and losers behave. Following Claim 4, a winner cannot benefit by raising her bid and a loser cannot benefit by lowering it, so we suppose that they never do this. Additionally, a loser who is actively engaged in the auction should raise her bid if it is beneficial. By Claim 5 we know that a loser can always raise her bid in a way that is weakly beneficial, so we suppose that a loser will always try to raise her bid.

- (A1). *A losing bidder will raise her bid in an effort to win; a loser will not lower her bid and a winner will not raise her bid.* Formally, if the current utility-target is π and i is a loser, then i must raise her bid at some point in the future unless she becomes a winner through the actions of other bidders.

Anecdotal evidence suggests that advertisers bidding in an ad auction generally expend substantial effort to launch advertising campaigns but are much slower to change things once they appear to work. Our second axiom generalizes this idea by supposing that winners (who, by definition, get the utility-target they request) view the outcome of the auction as a success while losers are unhappy with the results:

- (A2). *A bidder who is losing is more impatient than a bidder who is winning.* Formally, if the current utility-target is π and a set of bidders $L \subseteq [n]$ are losers, then the next time bids change it will necessarily be because some loser $i \in L$ raised her bid.

Our third axiom is analogous to (A1) but for winners — a winner who is actively engaged should lower her bid from time to time to see if she can win at a lower bid.

- (A3). *A winner will try lowering her effective bid to win at a lower price. Specifically, if a bidder is currently a winner, then she must lower her bid at some point in the future unless she becomes a loser through the action of another player.* Formally, if the current utility-targets are π and i is a winner, then i must lower her bid at some point in the future unless she becomes a loser through the actions of other bidders.

Our final axiom concerns the relative timing of events. Intuition and anecdotal evidence suggests that larger bidders who have more at stake tend to invest more heavily in active bidding strategies. This axiom roughly represents that intuition:

- (A4). *Between two losers, the bidder with the higher utility-target is more impatient.* Formally, if the current utility-targets are π and bidders i and j are both losers, then i will raise her bid before j if $\pi_i > \pi_j$.

These simple properties of bidder behavior imply the following convergence results. Proofs follow in Section 6.2 and Appendix A.

Theorem 6.1 *If losing bidders will only raise their effective bids (A1) and are more impatient than winning bidders (A2), the auction converges to the set of bids that are cooperatively envy-free (i.e. bids will be in C_ϵ).*

Theorem 6.2 *If winners try to lower their effective bids (A3) and losers try to raise but not lower their effective bids (A1), the auction converges to the set of bids that are non-cooperatively envy-free (i.e. bids will be in \bar{C}_ϵ).*

Combining Theorems 6.1 and 6.2 shows that bids will converge to the frontier of the CEF set. The strict Pareto frontier of this set is the set of CEF equilibrium bids.

Corollary 6.3 *If losing bidders will try raising their bids (A1), losers are less patient than winners (A2), and winners try lowering their bids (A3), the auction converges to the boundary between CEF and non-CEF bids (bids will be in the set $C_\epsilon \cup \bar{C}_\epsilon$).*

Finally, adding A4 induces convergence to a particular equilibrium:

Theorem 6.4 *If losing bidders will raise their effective bids (A1), winning bidders will try lowering their effective bids (A3), and the most impatient bidder is the losing bidder bidding for the highest utility-target (A2, A4), then bids will converge to the Egalitarian envy-free equilibrium.*

6.2 Convergence Proofs

In this section we give proofs of Theorems 6.1 and 6.2. Theorem 6.4 and omitted proofs may be found in Appendix A. Throughout this section, we assume that there is a single welfare optimal outcome for clarity of presentation.

Observation 3 *Under assumptions A1 and A2, a bidder will only lower her bid if all bidders are winners.*

Lemma 6.5 *If all bidders are winners under utility-targets π , then π is in the CEF set \mathcal{C} .*

Proof: If all bidders are winners, then we know that they are receiving precisely the utility-target they request when they bid π . Intuitively, this means that raising bids necessarily implies receiving less utility.

Formally, if bids are $b_i = (v_i, \pi_i)$ and the outcome of the auction is o^b , then we want to show that the CEF condition holds for any outcome o . Since all bidders are winners, we know $v_i(o^b) - b_i(o^b) = \pi_i$. Moreover, $v_i(o) - b_i(o) \leq \pi_i$ by definition, so $v_i(o) - b_i(o) \leq v_i(o^b) - b_i(o^b)$. Thus

$$\max\left((v_i(o) - b_i(o)) - (v_i(o^b) - b_i(o^b)), 0\right) = 0.$$

Since o^b is the outcome of the auction, we know $\sum_{i \in [n]} b_i(o^b) \geq \sum_{i \in [n]} b_i(o)$ for any outcome o . Thus, $0 \leq \sum_{i \in [n]} b_i(o^b) - b_i(o)$ and therefore

$$\sum_{i \in [n]} \max\left((v_i(o) - b_i(o)) - (v_i(o^b) - b_i(o^b)), 0\right) \leq \sum_{i \in [n]} b_i(o^b) - b_i(o)$$

as desired. ■

Since A1 and A2 imply that a player will only lower her bid from π if all bidders are winners, an important corollary is that a bidder will only lower her bid if the current utility-target vector is in the CEF set \mathcal{C} :

Corollary 6.6 *Under assumptions A1 and A2, if a player lowers her bid from π , then π is in the CEF set \mathcal{C} .*

A corollary of Claim 1 is that any set of CEF bids maximizes welfare, hence this implies that a player will only lower her bid if the welfare-optimal outcome is winning:

Corollary 6.7 *Under assumptions A1 and A2, a bidder will only lower her bid if a welfare-optimal outcome o^* is winning.*

Another useful fact about \mathcal{C} is that it is leftward-closed (the proof is in the appendix) and the natural corollary that $\overline{\mathcal{C}}$ is rightward-closed:

Lemma 6.8 *If π is in the CEF set \mathcal{C} , then $\pi - \Delta$ is in the CEF set \mathcal{C} for any $\pi \geq \Delta \geq 0$.*

Corollary 6.9 *If π is in the not-CEF set $\overline{\mathcal{C}}$, then $\pi + \Delta$ is in the not-CEF set $\overline{\mathcal{C}}$ for $\Delta \geq 0$.*

To prove that bids will converge, we first show that bids will not be stuck at arbitrarily low values:

Lemma 6.10 *Suppose the initial vector of utility-targets is π^0 . Under properties A1 and A2, the auction will always reach a configuration in which all bidders are winners and will do so within $\lceil \frac{1}{\epsilon} \pi^0 \rceil_1$ steps.*

We can now prove the our first theorem, that bids will be close to \mathcal{C} when A1 and A2 are satisfied.

Proof of Theorem 6.1: Lemma 6.10 implies that all bidders will be winners within a finite time. Once all bidders are winners, the only way bids will change is if someone lowers her bid. Thus, after a finite amount of time, we can conclude that either all bidders are winners or some bidder has lowered her bid.

Let π be the vector of utility-targets at any point after the first time all bidders are winners. If all bidders are still winners, then $\pi \in \mathcal{C}$ by Lemma 6.5. Otherwise, let i be the most recent player to lower her bid, increasing the utility-target vector from π' to $\pi'' = \pi' + \epsilon e_i$. We show that if i raises her bid again then the resulting utility-targets must be CEF regardless of how bids have changed since i 's raise.

By construction, players have only raised their bids since i lowered hers, so we can define $\Delta = \pi'' - \pi$ where $\Delta \geq 0$. Corollary 6.6 tells us that $\pi' \in \mathcal{C}$. If i raised her bid between π'' and π , then $\pi \leq \pi'$ and Lemma 6.8 tells us that $\pi \in \mathcal{C}$, so were done. Otherwise, we know $\pi'' \geq \Delta \geq 0$ and Lemma 6.8 tell us that $\pi' - \Delta \in \mathcal{C}$. Therefore $\pi = \pi' - \Delta + \epsilon e_j \in \mathcal{C}_\epsilon$. ■

To prove Theorem 6.2, we need a lemma similar to Lemma 6.10 showing that the auction will reach a bid vector that is CEF:

Lemma 6.11 *Under properties A1 and A3, as long as there is some outcome o and bidder j such that $v_j(o) > v_j(o^*)$, the auction will always reach a configuration that is not CEF when ϵ is sufficiently small. If there is no such outcome o and bidder j , then the auction may converge to $\pi_j = v_j(o^*)$ instead.*

Proof Sketch of Theorem 6.2: By Lemma 6.11, the auction will typically reach a utility-target vector in $\bar{\mathcal{C}}$. Our primary goal is to show that the auction will be in $\bar{\mathcal{C}}_\epsilon$ from that point onwards.

Suppose $\pi \in \mathcal{C}$. Let π' be the most recent utility-targets that were in $\bar{\mathcal{C}}$ and let $\pi'' \in \mathcal{C}$ be the utility-targets immediately after π' . Let i be the bidder who changed her bid between π' and π'' . Let o be an outcome that violated the CEF constraints at π' .

Consider a bidder j whose bid is higher at π than at π'' . Notice that j must be a loser to raise her bid, roughly $v_j(o^*) < \pi_j$. It then follows from the definition of the utility-target auction that $b_j(o^*)$ essentially did not change from π' to π . Moreover, $b_j(o)$ can only have increased from π' to π , so $b_j(o)$ increased more than $b_j(o^*)$ did. Similar logic leads to an analogous conclusion for bidders whose bids were lower at π'' than π , roughly giving

$$b_j(o) - b_j(o^*) \geq b'_j(o) - b'_j(o^*)$$

for each bidder $j \neq i$. For bidder i , a similar inequality holds accounting for the fact that i raised her bid from π' to π'' :

$$b_i(o) - b_i(o^*) \geq b'_i(o) - b'_i(o^*) - \epsilon .$$

Finally, given that the CEF constraint for o was violated at π' , these inequalities imply it must be nearly violated at π . ■

We have now shown that when losing bidders raise their effective bids and winning bidders lower their effective bids, bids remain close to the frontier of the CEF set \mathcal{C} . Adding in the specific

behavior that the first player to raise their bid will be the losing bidder with the highest utility-target results in convergence to one specific equilibrium: the egalitarian equilibrium (Theorem 6.4). The full proof is included in the appendix; we provide a sketch of it here.

Proof Sketch of Theorem 6.4:

Arrange bidders into levels L_1, \dots, L_k in increasing order of the utility each bidder gets at the egalitarian equilibrium.

For each level L_{i+1} , bids from all bidders in the level will converge close to the egalitarian equilibrium once the bids of lower level bidders are sufficiently close to their egalitarian bids.

Thus, beginning with the bidders who get the least utility in equilibrium, and working on up to the lucky bidders with the most utility, bids will converge close to the egalitarian outcome. ■

7 Conclusion and Open Questions

Pay-your-bid auctions — and utility-target auctions in particular — offer many advantages over incentive compatible mechanisms in terms of transparency and simplicity. Moreover, in many complex settings they even appear to generate more revenue.

Our work first shows that the bidding language is important in first-price auction design. In particular, it is both important and sufficient that bidders can compete in terms of their final utility. Also, a key feature in a repeated first-price auction is a pure-strategy equilibrium, something that GFP does not have [9]. This is a question of design: the existence of pure-strategy equilibria may be guaranteed through a carefully crafted bidding language (e.g. the utility-target auction) that can encode different per-click payments for different ad slots.

More significantly, when players compete on utility, our results show that robust performance guarantees may be derived using only simple axioms of bidder behavior that merely require knowledge of whether one is winning or losing. These results are powerful because they do not require an a priori assumption that the auction is in equilibrium or full information about others' bids.

Yet, reflection raises a concern about utility-target auctions: *why should bidders reveal their true valuation functions in a repeated auction?* We claimed that first-price auctions were better because the auctioneer could not cheat, but it would seem that quasi-truthfulness is just as dangerous. In fact, a quasi-truthful pay-your-bid auction is still strongly preferable to a standard second-price auction: even if the auctioneer knows a bidder's true valuation function, it cannot immediately increase the amount of money the bidder pays. By comparison, the auctioneer in a second-price auction might force a bidder to pay her full value in the second round by increasing the reserve price. The auctioneer is welcome to engage in a game of chicken or a "negotiation" with the bidder to see if she is willing to raise her bid, but the pay-your-bid property ensures that final approval still rests with the bidder.

In practice, systems may also be designed to encourage competition on the utility-target term and thereby recover stability. For example, Overture exacerbated the instability of the GFP auction by offering an API automating the sawtooth behavior. If an API were offered to compete on the utility-target term, bidders would likely use the API and stability would be restored, regardless of whether they were reporting their true valuation functions.

Issues of quasi-truthfulness aside, our work also raises questions about dynamic axioms of bidder behavior. Our axioms may be simple and natural, but strict adherence to them is clearly unrealistic.

In this vein, many interesting questions are open:

1. *How does the behavior of the auction change with small modifications to the axioms?* For example, we showed that bids would converge to the egalitarian equilibrium when the bidder with the most to gain raised first. Can we prove convergence to a different equilibrium by modifying players' delays?
2. *Do the performance guarantees still hold if axioms only hold probabilistically or on average?* It seems unlikely that bidder behavior always satisfies any particular set of axioms. How do the dynamic guarantees change when axioms only hold most of the time?
3. *What dynamic axioms do bidders actually obey?* An interesting experimental question is to determine what axioms are actually satisfied by bidder behavior. For example, could one experimentally measure bidders' delays and combine this with an answer to (1) to predict a particular equilibrium outcome?

References

- [1] Kenneth J. Arrow, H. D. Block, and Leonid Hurwicz. On the stability of the competitive equilibrium, ii. *Econometrica*, 27(1):82–109, 1959.
- [2] Itai Ashlagi, Mark Braverman, Avinatan Hassidim, Ron Lavi, and Moshe Tennenholtz. Position auctions with budgets: Existence and uniqueness. working paper, 2010.
- [3] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press, 2006.
- [4] Venkatesh Bala and Mukul Majumdar. Chaotic tatonnement. *Economic Theory*, 2(4):437–45, October 1992.
- [5] B. Douglas Bernheim and Michael D. Whinston. Menu Auctions, Resource Allocation, and Economic Influence. *The Quarterly Journal of Economics*, CI(1), 1986.
- [6] Richard Cole and Lisa Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 315–324, New York, NY, USA, 2008. ACM.
- [7] Robert Day and Paul Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3-4):393–407, July 2007.
- [8] Benjamin Edelman and Michael Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43(1):192–198, 2007.
- [9] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.
- [10] Benjamin Edelman and Michael Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions, 2010.

- [11] Richard Engelbrecht-Wiggans and Charles M. Kahn. Multi-unit pay-your-bid auctions with variable awards. *Games and Economic Behavior*, 23(1):25 – 42, 1998.
- [12] Eyal Even-dar, Yishay Mansour, and Uri Nadav. On the convergence of regret minimization dynamics in concave games. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 523–532, New York, NY, USA, 2009. ACM.
- [13] Glenn W. Harrison. Theory and misbehavior of first-price auctions. *The American Economic Review*, 79(4):pp. 749–762, 1989.
- [14] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [15] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions, 2012.
- [16] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce, EC '09*, pages 225–234, New York, NY, USA, 2009. ACM.
- [17] Darrell Hoy, Kamal Jain, and Christopher A. Wilkens. Coopetitive ad auctions. *arXiv*, abs/1209.0832, 2012.
- [18] Paul Klemperer. What really matters in auction design. *The Journal of Economic Perspectives*, 16(1):pp. 169–189, 2002.
- [19] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 869–886. SIAM, 2012.
- [20] David Lucking-Reiley. Vickrey auctions in practice: From nineteenth-century philately to twenty-first-century e-commerce. *Journal of Economic Perspectives*, 14(3):183–192, September 2000.
- [21] Eric Maskin and John Riley. Asymmetric auctions. *The Review of Economic Studies*, 67(3):pp. 413–438, 2000.
- [22] P. Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004.
- [23] Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):pp. 245–272, 2000.
- [24] Paul Milgrom and John Roberts. Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3(1):82 – 100, 1991.
- [25] Kevin Roberts. The characterization of implementable social choice rules. In *Aggregation and Revelation of Preferences*, J-J.Laffont (ed.), North Holland Publishing Company., 1979.
- [26] Tim Roughgarden and Éva Tardos. Do externalities degrade gsp’s efficiency? In *The Eighth Ad Auctions Workshop*, 2012.

- [27] Paul A. Samuelson. The stability of equilibrium: Comparative statics and dynamics. *Econometrica*, 9(2):97–120, 1941.
- [28] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [29] Leon Walras. *Elements of pure economics, or, The theory of social wealth / Leon Walras ; translated by William Jaffe*. Published for the American Economic Association and the Royal Economic Society by Allen and Unwin, London :, 1954.
- [30] Christopher A. Wilkens and Balasubramanian Sivan. Single-call mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, pages 946–963, New York, NY, USA, 2012. ACM.

A Selected Proofs

This appendix contains selected proofs that were previously omitted.

A.1 CEF Claims

Proof of Claim 1: We want to show that $\sum_{i \in [n]} v_i(o^*) \geq \sum_{i \in [n]} v_i(o)$ for any outcome o and CEF equilibrium o^* . The envy-freeness constraints give

$$\begin{aligned}
\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) &\leq \sum_{i \in [n]} b_i(o^*) - b_i(o) \\
\sum_{i \in [n]} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) &\leq \sum_{i \in [n]} b_i(o^*) - b_i(o) \\
\sum_{i \in [n]} v_i(o) - v_i(o^*) &\leq 0
\end{aligned}$$

as desired. ■

Proof of Claim 2: Define o^i as the outcome that maximizes the welfare of bidders except i :

$$o^i = \operatorname{argmax}_o \sum_{j \neq i} v_j(o) .$$

Thus, the VCG price of player i is $\sum_{j \neq i} v_j(o^i) - v_j(o^*)$.

Now, the envy-freeness constraints give

$$\begin{aligned}
b_i(o^*) &\geq \sum_j \max((v_j(o^i) - b_j(o^i)) - (v_j(o^*) - b_j(o^*)), 0) \\
&\quad + b_i(o^*) - \sum_j (b_j(o^*) - b_j(o^i)) \\
b_i(o^*) &\geq \sum_{j \neq i} (v_j(o^i) - b_j(o^i)) - (v_j(o^*) - b_j(o^*)) \\
&\quad - \sum_{j \neq i} (b_j(o^*) - b_j(o^i)) + b_i(o^i) \\
&\quad + \max((v_i(o^i) - b_i(o^i)) - (v_i(o^*) - b_i(o^*)), 0) \\
b_i(o^*) &\geq \sum_{j \neq i} (v_j(o^i) - v_j(o^*)) + b_i(o^i) \\
&\quad + \max((v_i(o^i) - b_i(o^i)) - (v_i(o^*) - b_i(o^*)), 0) \\
b_i(o^*) &\geq \sum_{j \neq i} (v_j(o^i) - v_j(o^*)) .
\end{aligned}$$

■

Proof of Claim 3: We want to show that

$$\sum_{i \in [n]} b_i(o^*) \geq \max_o \sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) .$$

For any outcome o , the envy-freeness constraints give

$$\begin{aligned}
\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) &\leq \sum_{i \in [n]} b_i(o^*) - b_i(o) \\
\sum_{i \in [n]} \max(v_i(o) - v_i(o^*) + b_i(o^*), b_i(o)) &\leq \sum_{i \in [n]} b_i(o^*) \\
\sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) &\leq \sum_{i \in [n]} \sum_{i \in [n]} b_i(o^*)
\end{aligned}$$

as desired. ■

A.2 Convergence Lemmas

Proof of Lemma 6.8: Let b be the bids at π and b^δ be the bids at $\pi - \delta$. Note that Claim 1 implies a welfare-optimal outcome o^* is winning at π .

First, suppose that all bidders for whom $\delta_i > 0$ are winners at π . In this case, $v_i(o^*) \geq \pi_i$ and so $v_i(o^*) \geq \pi_i - \delta_i$ and for any outcome o we get

$$\begin{aligned} \sum_{i \in [n]} b_i^\delta(o^*) &= \sum_{i \in [n]} b_i(o^*) + \delta_i \\ &\geq \sum_{i \in [n]} b_i(o) + \delta_i \\ &\geq \sum_{i \in [n]} b_i^\delta(o) , \end{aligned}$$

implying o^* is still winning at b_i^δ . Since $v_i(o^*) \geq \pi_i - \delta_i$, we can conclude that all bidders are winners, ergo $\pi - \delta \in \mathcal{C}$ by Lemma 6.5.

Now, suppose some bidders in π may be losers, but that the vector δ has the following property:

$$\delta_i \leq \max(\pi_i - v_i(o^*), 0) .$$

This condition says that only losers will raise their bids, and they will not raise them enough to affect $b_i(o^*)$.

Our goal is to show

$$\sum_{i \in [n]} b_i^\delta(o^*) - b_i^\delta(o) \geq \sum_{i \in [n]} \max\left((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\right) .$$

First, we see that $b_i^\delta(o) = b_i(o)$ as long as $v_i(o) \leq v_i(o^*)$. For any bidder i we have

$$b_i^\delta(o) = \max(v_i(o) - \pi_i - \delta_i, 0)$$

which can only be nonzero if $v_i(o) > \pi_i$. However, $b_i(o)$ can only change if $\delta_i > 0$, which requires $\pi_i > v_i(o^*)$ and thus $v_i(o) > \pi_i > v_i(o^*)$. By construction, this also holds for $\pi_i - \delta_i$:

$$v_i(o) > \pi_i - \delta_i \geq v_i(o^*) .$$

Now, when $v_i(o) > \pi_i - \delta_i \geq v_i(o^*)$, we have

$$\begin{aligned} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)) &= \min(v_i(o), \pi_i - \delta_i) - \min(v_i(o^*), \pi_i - \delta_i) \\ &= \pi_i - \delta_i - v_i(o^*) \\ &\geq 0 . \end{aligned}$$

Importantly, if $\Delta(o)$ is the set of bidders for which $b_i^\delta(o) \neq b_i(o)$, we may conclude that

$$\begin{aligned} \sum_{i \in [n]} \max\left((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\right) &= \\ &= \sum_{i \notin \Delta(o)} \max\left((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\right) \\ &\quad + \sum_{i \in \Delta(o)} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)) \end{aligned}$$

and likewise

$$\begin{aligned}
& \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) = \\
& = \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) \\
& \quad + \sum_{i \in \Delta(o)} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) .
\end{aligned}$$

The desired CEF condition quickly follows, using the fact that bidders $i \notin \Delta(o)$ did not change their bids:

$$\begin{aligned}
& \sum_{i \in [n]} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) = \\
& = \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) + \sum_{i \in \Delta(o)} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)) \\
& = \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) + \sum_{i \in \Delta(o)} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) \\
& \quad + \sum_{i \in \Delta(o)} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*)) \\
& = \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) + \sum_{i \in [n]} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*)) \\
& \leq \sum_{i \in [n]} (b_i(o^*) - b_i(o)) + \sum_{i \in [n]} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*)) \\
& = \sum_{i \in [n]} b_i^\delta(o^*) - b_i^\delta(o)
\end{aligned}$$

as desired.

Finally, for general δ , split it as $\delta = \delta^1 + \delta^2$ where

$$\delta_i^1 = \min(\delta_i, \pi_i - v_i(o^*)) .$$

The vector δ^1 satisfies the condition $\delta_i \leq \max(\pi_i - v_i(o^*), 0)$, so $\pi - \delta^1 \in \mathcal{C}$. Moreover, all bidders are winners in $\pi - \delta^1$, so

$$s - \delta^1 - \delta^2 = s - \delta \in \mathcal{C}$$

as desired. ■

Proof of Lemma 6.10: Properties A1 and A3 imply that a bid will only be lowered if there are no losers. Thus, bids will only be raised (utility-targets decreased) until all bidders are simultaneously winners. Since any bidder i is always a winner when bidding $\pi_i = 0$ and bidders never decrease their utility-targets when they are winners (A1), utility-target can be decreased at most

$\lceil \frac{1}{\epsilon} s \rceil_1$ times before all bidders are winners. Moreover, since losers will always try to decrease their utility-targets (A1), the auction will never stall in a configuration where some bidder is a loser. ■

Proof of Lemma 6.11: If the auction reaches a vector π that induces an outcome $o \neq o^*$, then $\pi \in \overline{\mathcal{C}}$ and we are done. Thus, it remains to show that an auction will reach a vector $\pi \in \overline{\mathcal{C}}$ even if the outcome is always o^* .

Consider a bidder j . By A1 we know that j will only decrease π_j if she is a loser and increase π_j if she is a winner. By A3 we can conclude that j will eventually decrease her bid until $\pi_j \geq v_j(o^*) - \epsilon$, implying $b_j(o^*) \leq \epsilon$. Thus,

$$\sum_{i \in [n]} b_i(o^*) - b_i(o) \leq n\epsilon .$$

Now, as long as there is some outcome o and bidder j such that $v_j(o) > v_j(o^*)$, when ϵ is sufficiently small it will be the case that

$$n\epsilon < \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0)$$

Unfortunately, this implies

$$\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) > \sum_{i \in [n]} b_i(o^*) - b_i(o) ,$$

and therefore $\pi \in \overline{\mathcal{C}}$.

If there is no outcome o and bidder j such that $v_j(o) > v_j(o^*)$, then by similar logic bidders will increase their utility-targets until precisely $\pi_j = v_j(o^*)$ (as a result of our restriction that bidders always bid $\pi_j \leq \sup_o v_j(o^*)$). ■

Theorem 6.2 (Restatement). *If winners try to lower their effective bids (A3) and losers try raising but not lowering (A1), bids will eventually remain close to the boundary of the set of CEF bids or entirely outside it.*

Proof of Theorem 6.2: By Lemma 6.11, the auction will eventually reach a utility-target vector in $\overline{\mathcal{C}}$ or the degenerate case where nobody is paying anything and o^* is winning. In the degenerate case, bids converge to a point on the boundary of \mathcal{C} , so the theorem is true. For the standard case, we show that $\pi \in \overline{\mathcal{C}}_\epsilon$ from the first time a bid in $\overline{\mathcal{C}}$ is reached.

If $\pi \in \overline{\mathcal{C}}$, we are done, so suppose $\pi \in \mathcal{C}$. Let π' be the most recent utility-targets that were in $\overline{\mathcal{C}}$ and let $\pi'' \in \mathcal{C}$ be the utility-targets immediately after π' . Let i be the bidder who changed her bid between π' and π'' . Corollary 6.9 implies that i must have raised her bid between π' and π'' .

First, suppose that the outcome changed from o' to o'' when i raised her bid. Since o^* must be the outcome of any CEF bid, we know that $o'' = o^*$ and that the outcome does not change again before bids reach π . Define the utility-target vector \tilde{s} with associated bids \tilde{b} as follows:

$$\tilde{\pi}_j = \begin{cases} \min(\pi_i, \pi'_i) & j = i \\ \pi_j - \epsilon & \pi_j > \pi''_j \\ \pi_j + \epsilon & \pi_j < \pi'_j \\ \pi_j & \text{otherwise.} \end{cases}$$

Let $\delta_j = \tilde{\pi}_j - \pi_j''$. We argue later that i will not increase her utility-target from $\pi_i'' = \pi_i' + \epsilon$, so $|\tilde{\pi}_j - \pi_j| \leq \epsilon$ for all j . Thus, it is sufficient to show that $\tilde{s} \in \bar{\mathcal{C}}$.

Consider a bidder $j \neq i$ and suppose $\pi_j > \pi_j''$. By definition, we get a simple bound on j 's bid for o' :

$$\tilde{b}_j(o') \geq b'_j(o') - \delta_j .$$

We also know that j lowered her bid at some point between π'' and π . Since j would only increase her utility-target if she were a winner, she must have been a winner at some value $\geq \pi_j - \epsilon = \tilde{\pi}_j$. Thus, $\tilde{\pi}_j = \pi_j - \epsilon \leq v_j(o^*)$. We can thus upper-bound her bid for o^* :

$$\tilde{b}_j(o^*) \leq b'_j(o^*) - \delta_j .$$

Combining these two bounds and noting that $b_j'' = b'_j$ for $j \neq i$ gives

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*) .$$

For bidders $j \neq i$ with $\pi_j < \pi_j''$, analogous reasoning based the fact that j must have been a loser to decrease her utility-target gives

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*) .$$

For bidders $j \neq i$ with $\pi_j = \pi_j''$, we trivially have

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) = b'_j(o') - b'_j(o^*) ,$$

so it remains to consider bidder i .

For bidder i , we know that decreasing her utility-target from π_i' to π_i'' increased her bid for o^* more than it increased her bid for o' . This implies $v_i(o') < v_i(o^*)$ and $\pi_i'' < v_i(o^*)$. Consequently, i is a winner with π_i'' at o^* and will not decrease her utility-target further. Firstly, this implies that $|\tilde{\pi}_i - \pi_i| \leq \epsilon$. First, suppose $\pi_i > \pi_i''$. In this case, $\pi_i \geq \pi_i'$, and since $v_i(o^*) > v_i(o')$ we have

$$\tilde{b}_i(o') - \tilde{b}_i(o^*) \geq b'_i(o') - b'_i(o^*) .$$

Otherwise, i does not change her bid from π'' to π , so $\tilde{\pi}_i = \pi_i'$ and therefore $\tilde{b}_i = b'_i$

Thus, for any bidder j we have

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*) ,$$

and thus

$$\sum_{j \in [n]} \tilde{b}_j(o') - \sum_{j \in [n]} \tilde{b}_j(o^*) \geq \sum_{j \in [n]} b'_j(o') - \sum_{j \in [n]} b'_j(o^*) .$$

Since o' was winning at b'_j , this implies o^* cannot be winning under $\tilde{\pi}$, and therefore $\tilde{\pi} \in \bar{\mathcal{C}}$. By construction, $|\tilde{\pi}_j - \pi_j| \leq \epsilon$, so this implies $\pi \in \bar{\mathcal{C}}_\epsilon$.

So far, we showed that $\pi \in \bar{\mathcal{C}}_\epsilon$ as long as the outcome changed when i raised her bid. In the case where the outcome was already $o' = o^*$, we want to analyze the CEF constraints directly. Since $\pi' \in \bar{\mathcal{C}}$, there is some outcome o^w for which the CEF constraints are violated, i.e.

$$\sum_{i \in [n]} b'_i(o^*) - b'_i(o^w) < \sum_{i \in [n]} \max((v_i(o^w) - b'_i(o^w)) - (v_i(o^*) - b'_i(o^*)), 0) .$$

Observing that the outcome does not change from π' to π , the logic from the case where $o' \neq o^*$ gives

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*)$$

for any bidder j . It immediately follows that

$$\sum_{i \in [n]} \tilde{b}_i(o^*) - \tilde{b}_i(o^w) < \sum_{i \in [n]} \max \left((v_i(o^w) - \tilde{b}_i(o^w)) - (v_i(o^*) - \tilde{b}_i(o^*)), 0 \right) ,$$

and so $\tilde{\pi} \in \overline{\mathcal{C}}$ and $\pi \in \overline{\mathcal{C}}_\epsilon$. ■

A.3 Convergence to the Egalitarian Equilibrium

Theorem 6.4 (Restatement). *If losing bidders will raise their effective bids (A1), winning bidders will try lowering their effective bids (A3), and the most impatient bidder is the losing bidder bidding for the highest utility (A2, A4), then bids will converge to the Egalitarian envy-free equilibrium.*

Proof of Theorem 6.4: The proof will proceed as follows. We first categorize bidders into levels based on their utility in the egalitarian outcome. We define upper and lower bounds on utility-targets as multiples of ϵ , the amount by which players change their bids. Then, we show that if for a given bidder j , the bid of every lower-utility bidder has converged to within their bounds, the bid of j will also converge to within her bounds - first to at least her lower bound (Lemma A.1), and then to at most her upper bound (Lemma A.2). Combining these via induction gives our final result that the bids of all players converge near their egalitarian outcome.

Let o^* be the egalitarian outcome; let π_i^* be the corresponding utility-target of bidder i . Let $\mathbf{B}_b(o) = \sum_{i \in [n]} b_i(o)$ be the total bid for a given outcome o . Let $\mathbf{B}_X(o) = \sum_{j \in L_X} b_j(o)$, and $\mathbf{B}_X^*(o)$ be similarly defined.

First, consider all utility-targets in the egalitarian equilibrium; let z_i be the i th smallest (distinct) utility-target. Let L_i be the set of all players with a utility-target of z_i in the egalitarian equilibrium. We will use $L(j)$ to denote the level of a bidder j .

We will show convergence by showing that there exist functions $b^-(i)$ and $b^+(i)$ s.t. for any $j \in L_i$, utility-targets converge into and remain in the interval $[\pi_j^* - \epsilon b^-(i), \pi_j^* + \epsilon b^+(i)]$.

Bidding Bounds. We now precisely define the bounding functions $b^-(\cdot)$ and $b^+(\cdot)$.

Definition 11

$$b^-(i) = 2^{2|L_{<i}|} \tag{1}$$

$$b^+(i) = 2^{2|L_{<i}| + |L_i|} \tag{2}$$

These bounds are given specifically so that for any level k , the sum over upper bounds in lower levels is at most the lower bound in level k , and the sum over all lower bounds for lower (or equal) levels is at most the upper bound for level k . Intuitively, we are saying that lower-level bidders cannot over or under bid enough to make up for bidders in level k .

Claim 6

$$b^-(k) > \sum_{i=0}^{k-1} |L_i| b^+(i) \quad (3)$$

$$b^+(k) > \sum_{i=0}^k |L_i| b^-(i) \quad (4)$$

We omit the proof; it follows from manipulation of exponential sums.

Witness outcomes. Recall from Algorithm 2 that utility-targets for any given player are raised until the CEF constraint for some outcome o is violated. These outcomes have an important role to play in the egalitarian equilibrium — they are the reason that a bidder cannot achieve any more utility. We will call them *witness* outcomes.

Three properties of these witness outcomes are important for us. First, bidder i values the witness at less than her egalitarian bid, hence she would be ‘losing’ if it was chosen above the egalitarian winning ad; that all bidders with higher utility value it at at least their utility-target; and that with the final winning bids, the total bid of each is tied. We define *witness* outcomes precisely as follows:

Definition 12 *Outcome o^w is a witness outcome for bidder j at the egalitarian utility-targets π^* if its total bid is tied with the egalitarian outcome, j asking for more utility at the egalitarian equilibrium results in a higher total bid for o^w than for the optimal egalitarian ad and j is the highest-utility bidder to lose if o^w wins over o^* .*

Recall the intuition behind these outcomes: they are the reason that a bidder cannot achieve more utility at the egalitarian equilibrium. If there is no witness for a bidder who must pay something, then the bidder could ask for more utility, and higher utility bidders could effectively ‘pick up the slack’, resulting in a more egalitarian outcome.

Claim 7 *At the egalitarian equilibrium, every bidder j s.t. $\pi_i^* < v_i(o^*)$ has at least one witness.*

Proof: We will prove via contradiction. Assume at the CEF egalitarian outcome o^* bidder j has no witness. Now, let bidder j increase her utility-target by a small enough $\epsilon > 0$, that only outcomes that were previously tied with o^* win over o^* . For each of these outcomes, there must be a higher utility bidder than j who does not win with the outcome; otherwise it would be a witness for j . Decrease the utility-targets of the highest utility bidder not in each of these outcomes by ϵ . At this point, all outcomes will be tied again — and we can have the optimal outcome win the tiebreaker via having a higher utility, or assume that one player will decrease, then raise their utility-target to ensure that it was the previous outcome to win. These bids will be CEF, and will be more egalitarian than o^* , as bidder j achieved more utility, and only higher utility bidders achieved less utility. ■

Another important property will be that each outcome is only a witness for bidders of a single level:

Claim 8 *An outcome is only a witness outcome for bidders of a single level.*

This really follows from the definition — bidders in different levels cannot both be the highest utility bidder to not win with an outcome. More intuitively though, if players of different levels were both not in an outcome, and the lower utility bidder had no other witness outcome, then a more egalitarian outcome would involve increasing his utility-target, and decreasing the utility-target of the higher utility bidder.

Bidding convergence. We now present the core of our convergence result. This convergence is a two step process for bidders in a given level; after the utility bids of all lower level bidders have converged within their bounds, convergence in the given level to at least the lower bound takes place first, and then bids in the given level will converge to below their upper bound.

Lemma A.1 *Under assumptions A1, A2, A3 and A4, the utility-target of each bidder j in level L_i will converge to at least their lower bounds, $\pi_j^* - \epsilon \cdot b^-(i)$ if for every bidder j' in level L_{i-} s.t. $i^- < i$, $\pi_{j'} \leq \pi_{j'}^* + \epsilon \cdot b^+(i^-)$.*

Proof: Our argument consists of two parts: first, that if a bidder is bidding for utility at or below her lower bound then she will never reduce her utility-target further. Second, she will eventually try raising her bid (by A3). These two combined will lead to her eventually raising her bid to at least the lower bound.

Claim 9 *Under the assumptions of Lemma A.1, no bidder j in level L_i with a utility-target of $\pi_j \leq \pi_j^* - \epsilon \cdot b^-(i)$ will lower her utility-target.*

We will prove via contradiction. Assume for bidder j that with a utility-target of $\pi_j \leq \pi_j^* - \epsilon \cdot b^-(i)$, she wishes to lower her utility-target further. Let o be the winning outcome with bids π . As i will only lower her bid if she is losing (A1), $\pi_j > v_j(o)$. We will now try to derive the contradiction that the total bid for the optimal outcome is at least the total effective bid for o ($\mathbf{B}(o^*) > \mathbf{B}(o)$), hence she must win and would not care to lower her utility-target.

For j to decrease her utility-target, by A4 she must be the highest utility bidder who is losing such that $\pi_j > v_j(o)$. By our bound, we know that for every lower utility bidder j' in level L_{i-} , $\pi_{j'} \leq \pi_{j'}^* + \epsilon b^+(i^-)$. Since o^* is the optimal winning outcome and o the currently winning outcome, $\mathbf{B}^*(o^*) \geq \mathbf{B}^*(o)$ and $\mathbf{B}(o^*) \leq \mathbf{B}(o)$.

In the egalitarian outcome, every bidder j receives the utility she bids for; hence $b_j^*(o^*) = v_j(o^*) - \pi_j^*$. By our assumption on the utility-target bounds, for all bidders $j' \in L_{<i+1}$, $b_{j'}^*(o^*) - b_{j'}(o^*) \leq \epsilon b^+(L(j'))$.

Consider a bidder j' in a lower level than j and first, is requesting more utility relative to the egalitarian outcome, specifically that $\pi_{j'} > \pi_{j'}^*$. Hence, we will have $0 \geq b_{j'}(o^*) - b_{j'}^*(o^*) \geq -(\pi_{j'} - \pi_{j'}^*)$ and $0 \geq b_{j'}(o) - b_{j'}^*(o) \geq -(\pi_{j'} - \pi_{j'}^*)$. Hence,

$$(b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o) - b_{j'}^*(o)) \geq -(\pi_{j'} - \pi_{j'}^*) \quad (5)$$

$$\geq -b^+(L(j')). \quad (6)$$

Consider the case that $\pi_{j'} \leq \pi_{j'}^*$, that j' is requesting less utility than in the egalitarian outcome. Then $b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*) \geq 0$, and $b_{j'}(o) - b_{j'}^*(o) \leq -(\pi_{j'} - \pi_{j'}^*)$. Hence,

$$(b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o) - b_{j'}^*(o)) \geq 0. \quad (7)$$

Summing over all lower-level bidders via Equations (6) and (7) gives $(\mathbf{B}_{<i}(o^*) - \mathbf{B}_{<i}^*(o^*)) - (\mathbf{B}_{<i}(o) - \mathbf{B}_{<i}^*(o)) \geq -\sum_{i' < L(j)} b^+(i')$ and hence by Claim 6,

$$(\mathbf{B}_{<i}(o^*) - \mathbf{B}_{<i}^*(o^*)) - (\mathbf{B}_{<i}(o) - \mathbf{B}_{<i}^*(o)) > -b^-(i). \quad (8)$$

Now, consider a bidder j' in the same or a higher level than j . If j' is overbidding and not winning in outcome o with bids b , then she would have decreased her utility-target faster than j . She could however be overbidding and winning in o ; in which case the decrease in bids for o^* must be bounded by the decrease for o , hence: $(b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o) - b_{j'}^*(o)) \geq 0$. If she is requesting less utility, o^* will see the full increase in bid while o may not. Denote the total bid of all bidders aside from j in the same or higher level as j as $\mathbf{B}_{\geq i \setminus j}(o)$. Then, summing over all such bidders gives

$$\mathbf{B}_{\geq i \setminus j}(o^*) - \mathbf{B}_{\geq i \setminus j}^*(o^*) - (\mathbf{B}_{\geq i \setminus j}(o) - \mathbf{B}_{\geq i \setminus j}^*(o)) \geq 0. \quad (9)$$

Our original assumption on j gives $(b_j(o^*) - b_j^*(o^*)) - (b_j(o) - b_j^*(o)) \leq b^-(i)$. Now, taking the sum over this and equations (8) and (9) gives $(\mathbf{B}(o^*) - \mathbf{B}^*(o^*)) - (\mathbf{B}(o) - \mathbf{B}^*(o)) > -b^+(L(j)) + b^+(L(j)) = 0$. By our assumption that o^* is the egalitarian winning outcome, we have $\mathbf{B}^*(o^*) - \mathbf{B}^*(o) \geq 0$. Adding these yields

$$\mathbf{B}(o^*) - \mathbf{B}(o) > 0. \quad (10)$$

This is in violation of our assumption that o wins with bids b . Hence, no such bidder j can ever wish to lower her utility-target past the lower bound when all lower-level agents have bids within their upper bounds. By Assumption A3, she will eventually try and lower her bid when winning, hence her bid will converge above her lower bound. \blacksquare

Lemma A.2 *Under assumptions A1, A2, A3 and A4, the utility-target π_j of each bidder j in level L_i will converge to at most the upper bound, $\pi_j^* + \epsilon \cdot b^+(i)$ if for every bidder j' in level L_{i-} s.t. $i^- \leq i$, $\pi_{j'} \geq s_{j'}^* - \epsilon \cdot b^-(i^-)$.*

Proof:

By Assumptions A1, A2 and Observation 3, a bidder will only request more utility from a set of bids b with winning outcome o if all other bidders are winning with bids b , and by Lemma 6.5, b must be CEF.

Our proof will proceed by showing that in any such o , $\pi_j < \pi_j^* + \epsilon \cdot b^+(i)$, and hence her utility-target must stay below $\pi_j^* + \epsilon \cdot b^+(i)$ in winning outcomes. Furthermore, by Theorem 6.1 bids will become CEF; hence i will be forced to decrease her utility-target.

By Claim 7, there is a witness outcome o^w which includes every bidder j' in a strictly higher level i^+ than j . We will now show that if all other players are winning with the egalitarian winning outcome, then j 's utility-target must be below her upper bound, otherwise the witness outcome o^w would win over o^* .

By Definition 12, $\mathbf{B}^*(o^w) = \mathbf{B}^*(o^*)$. Consider the quantity $\mathbf{B}(o^*) - \mathbf{B}^*(o^*)$, and break it into sums over bidders in levels at or below bidder j , j and bidders in levels above j :

$$\begin{aligned} (\mathbf{B}(o^*) - \mathbf{B}^*(o^*)) &= (\mathbf{B}_{\leq i \setminus j}(o^*) - \mathbf{B}_{\leq i \setminus j}^*(o^*)) \\ &\quad + (b_j(o^*) - b_j^*(o^*)) \\ &\quad + (\mathbf{B}_{>i}(o^*) - \mathbf{B}_{>i}^*(o^*)) \end{aligned}$$

We will now proceed by separately considering bidders in higher and lower levels than bidder j . We will bound the change in bids from each, and see that there is no way for bidder j to ask for utility above her upper bound and still be in the winning outcome.

Higher-level bidders. By properties of witness sets, any such bidder j' must be winning in the witness outcome at both the egalitarian bids and the current bids, hence $b_{j'}(o^w) - b_{j'}^*(o^w) = -(\pi_{j'} - \pi_{j'}^*)$. Since we know that at the egalitarian bids, such a bidder must be winning in the egalitarian outcome, $b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*) \leq b_{j'}(o^w) - b_{j'}^*(o^w)$. Summing over all such bidders yields

$$(\mathbf{B}_{>i}(o^*) - \mathbf{B}_{>i}^*(o^*)) - (\mathbf{B}_{>i}(o^w) - \mathbf{B}_{>i}^*(o^w)) \leq 0 \quad (11)$$

Lower-level bidders. By our initial assumption that bidding has converged above lower bounds for these bidders, for any bidder j' in $L_{\leq i}$, $\pi_{j'} \geq \pi_{j'}^* - \epsilon b^-(L(j'))$, and hence $b_{j'}(o^*) \leq b_{j'}^*(o^*) + \epsilon b^-(L(j'))$ and $b_{j'}(o^w) \leq b_{j'}^*(o^w) + \epsilon b^-(L(j'))$.

Recall that all bids $b_{j'}^*(o^*)$ and $b_{j'}(o^*)$ are winning by assumption — since o^* is the egalitarian outcome, and no player wishes to decrease their utility-target in the current bids. If for some bidder j' , $v_{j'}(o^w) \geq v_{j'}(o^*)$, then $(b_{j'}(o^w) - b_{j'}^*(o^w)) = (b_{j'}(o^*) - b_{j'}^*(o^*)) = -(\pi_{j'} - \pi_{j'}^*)$.

Consider then the case that $v_{j'}(o^w) < v_{j'}(o^*)$; that is, that j' values the witness o less than the egalitarian outcome. We will consider two cases: that her utility-target is lower or higher than her egalitarian utility-target respectively.

$(\pi_{j'} < \pi_{j'}^*)$ If the bidder j' bids for less utility than in the egalitarian outcome, then that increase in effective bid will be bounded by the increase in the bid for the egalitarian outcome. That is, we have $b_{j'}(o^w) - b_{j'}^*(o^w) = \max(v_{j'}(o^w), \pi_{j'}) - \pi_{j'} - \max(v_{j'}(o^w), \pi_{j'}^*) + \pi_{j'}^*$, and hence $b_{j'}(o^w) - b_{j'}^*(o^w) = -(\pi_{j'} - \pi_{j'}^*) + (\max(v_{j'}(o^w), \pi_{j'}) - \max(v_{j'}(o^w), \pi_{j'}^*))$. As $b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*)$, we then have:

$$0 \leq b_{j'}(o^w) - b_{j'}^*(o^w) \leq b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*). \quad (12)$$

Furthermore, since $\pi_{j'} \geq \pi_{j'}^* - \epsilon b^-(L(j'))$ by assumption, we have:

$$0 \leq b_{j'}(o^w) - b_{j'}^*(o^w) \leq b_{j'}(o^*) - b_{j'}^*(o^*) \leq \epsilon b^-(L(j')) \quad (13)$$

and

$$0 \leq (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w)) \leq \epsilon b^-(L(j')). \quad (14)$$

$(\pi_{j'} \geq \pi_{j'}^*)$ If bidder j' instead is bidding for at least as much utility as in the egalitarian outcome, the decrease in total bid is bounded by the change in bids for the egalitarian outcome, hence the change in utility-targets will be between $b_{j'}(o^*) - b_{j'}^*(o^*) = \pi_{j'}^* - \pi_{j'}$ and 0. Hence,

$$b_{j'}(o^*) - b_{j'}^*(o^*) \leq b_{j'}(o^w) - b_{j'}^*(o^w) \leq 0 \quad (15)$$

and

$$-(\pi_{j'} - \pi_{j'}^*) \leq (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w)) \leq 0. \quad (16)$$

We now have upper bounds on $-(\pi_{j'} - \pi_{j'}^*) \leq (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w))$ for all lower-level bidders. Taking the sum across all members of $L_{\leq i}$ via Equations (14) and (16) gives:

$$\sum_{j' \in L_{\leq i}} (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w)) \leq \sum_{j' \in L_{\leq i}} \epsilon b^-(L(j')) \quad (17)$$

Rearranging and noting that $\sum_{j' \in L_{\leq i}} \epsilon b^-(L(j')) < b^+(i)$ by Claim 6 gives

$$(\mathbf{B}_{\leq i \setminus j}(o^*) - \mathbf{B}_{\leq i \setminus j}^*(o^*)) - (\mathbf{B}_{\leq i \setminus j}(o^w) - \mathbf{B}_{\leq i \setminus j}^*(o^w)) < \epsilon b^+(i). \quad (18)$$

Summing over equations (18) and (11) gives us:

$$\begin{aligned} & (\mathbf{B}_{\leq i \setminus j}(o^*) - \mathbf{B}_{\leq i \setminus j}^*(o^*)) - (\mathbf{B}_{\leq i \setminus j}(o^w) - \mathbf{B}_{\leq i \setminus j}^*(o^w)) \\ & + (\mathbf{B}_{> i}(o^*) - \mathbf{B}_{> i}^*(o^*)) - (\mathbf{B}_{> i}(o^w) - \mathbf{B}_{> i}^*(o^w)) < \epsilon b^+(i). \end{aligned} \quad (19)$$

By assumption, $(b_j(o^*) - b_j^*(o^*)) = -(\pi_j - \pi_j^*) \leq -\epsilon b^+(i)$ and $b_j(o^w) = b_j^*(o^w) = 0$. Thus, $(b_j(o^*) - b_j^*(o^*)) - (b_j(o^w) - b_j^*(o^w)) \leq -\epsilon b^+(i)$. Adding this to (19) gives:

$$(\mathbf{B}(o^*) - \mathbf{B}^*(o^*)) - (\mathbf{B}(o^w) - \mathbf{B}^*(o^w)) < \epsilon b^+(i) - \epsilon b^+(i) = 0 \quad (20)$$

By our initial assumption that o^w is a witness outcome, $\mathbf{B}^*(o^*) - \mathbf{B}^*(o^w) = 0$. Adding this to the above equation yields

$$\mathbf{B}(o^*) < \mathbf{B}(o^w) \quad (21)$$

This contradicts our assumption that o^* is a winning set with bids $b(\cdot)$. Hence, j will be forced to decrease her utility-target to at most $\pi_j^* + \epsilon b^+(i)$ before the egalitarian winning set o^* is winning again. ■

Combining Lemma A.1 and Lemma A.2 gives us convergence of each bidder in each level i to within their bounds as soon as lower level bidders have all converged. It follows then from straightforward induction on levels that all bids converge to within their bounds. ■